arXiv:1703.09865v1 [cs.NE] 29 Mar 2017

Experience-based Optimization: A Coevolutionary Approach

Shengcai Liu[a], Ke Tang[a,*], Xin Yao[b]

[a] *USTC-Birmingham Joint Research Institute in Intelligent Computation and Its Applications (UBRI), School of Computer Science and Technology, University of Science and Technology of China, Hefei, Anhui, China*
[b] *Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen, China*

Abstract

This paper studies improving solvers based on their past solving experiences, and focuses on improving solvers by offline training. Specifically, the key issues of offline training methods are discussed, and research belonging to this category but from different areas are reviewed in a unified framework. Exsiting training methods generally adopt a two-stage strategy in which selecting the training instances and training instances are treated in two independent phases. This paper proposes a new training method, dubbed LiangYi, which addresses these two issues simultaneously. LiangYi includes a training module for a population-based solver and an instance sampling module for updating the training instances. The idea behind LiangYi is to promote the population-based solver by training it (with the training module) to improve its performance on those instances (discovered by the sampling module) on which it performs badly, while keeping the good performances obtained by it on previous instances. An instantiation of LiangYi on the Travelling Salesman Problem is also proposed. Empirical results on a huge testing set containing 10000 instances showed LiangYi could train solvers that perform significantly better than the solvers trained by other state-of-the-art training method. Moreover, empirical investigation of the be-

*Corresponding author
Email addresses: liuscyyf@mail.ustc.edu.cn (Shengcai Liu), ketang@ustc.edu.cn (Ke Tang), xiny@sustc.edu.cn (Xin Yao)

haviours of LiangYi confirmed it was able to continuously improve the solver through training.

Keywords: experience, combinatorial optimization, offline training, coevolution, dynamic sampling, parallel solvers

---

1. Introduction

Hard optimization problems (e.g., NP-hard problems) are ubiquitous in AI research and real-world applications. To tackle them, numerous problem solvers have been proposed over the last few decades [**?**]. In general, such a solver is designed for a certain problem domain rather than a single instance, because when used in practice, it usually needs to solve many different instances belonging to that domain. Besides, many of such problem solvers are heuristic methods, the performance (e.g., time complexity required to obtain the optimal solution) of which can hardly be rigorously proved. As a result, the development of these solvers typically involves repeatedly testing it against a number of problem instances and adjusting it based on the testing results [**?**].

Given that both the design and the applications of a solver would involve many problem instances, a natural question is whether a solver could leverage on the experience acquired from solving previous problem instances to grow/enhance its capacity in solving new coming problem instances. For the sake of brevity, a solver whose performance can improve as it solves more and more problem instances is termed as an Experience-based Optimizer (EBO) in this work.

The intuitions behind EBO are two-fold. First, any human expert in a specific domain starts as a novice and his/her path to an expert mainly relies on the gradual accumulation of problem-solving experience in this domain. Second, exploiting past experience to facilitate the solving of new problems, from a more technical point of view, concerns the generalization of past experience, which lies in the heart of AI research, particularly the machine learning sub-area. The past few decades have witnessed great progress on this issue, while most successes

were achieved on building a learner that can correctly map an input signal (e.g., an image) to a predefined output (e.g., a label). It is interesting to ask whether similar idea could be developed to encompass more complex problem such as NP-hard optimization problems, which may introduce new challenges as the desired output will no longer be a label (or other types of variables), but a solution to the optimization problem.

EBO could offer three advantages in practice. First, it enables an automated process analogous to life-long learning of we human beings and thus alleviate the tedious step-by-step fine-tuning or upgrade work that is now mostly done by human experts. Second, as an EBO can improve its performance automatically, it would be able to better exploit nowadays high-performance computing facilities to generate and test much more problem instances than a domain expert can do manually, such that the risk of over-tuning the solver to a small set of problem instances can be reduced. Finally, the underlying properties of real-world hard optimization problem instances, even if they are from exactly the same problem class, may change over time. Since EBO dynamically updates itself when solving more and more problem instances, it would better fit the changing world.

Analogous to many machine learning techniques, an EBO may run in two manners, i.e., the offline and online modes. For the offline mode, a set of problem instances is fed to the EBO at a time and the solver is updated after collecting the solutions it obtained on all the instances. For the online mode, problem instances are fed to EBO one at a time and the solver is updated immediately after solving an instance. Being conceptually different, the offline mode might play a more important role in EBO since a set of training instances are usually available when designing a solver. On comparison, the online mode is more likely to occur after the EBO is deployed in a real-world application. Even in this case, offline mode could also be adopted since EBO simply needs to postpone its update until collecting more instances. For these considerations, we focus on the offline mode as the first step to investigate EBO. Specifically, this work consists of three main parts, as summarized below:

(1) An overview on the key issues of EBO: In the literature, there have been several pioneer attempts to design mechanisms that enhance a solver based on past experience. For examples, Transfer methods [**? ? ?**], as the name implies, transfer the useful information extracted from solved instances to unsolved problem instances. Automatic algorithm configuration [**? ?**], Portfolio-based algorithm selection [**? ? ?**], and Automatic portfolio construction [**? ? ? ?**] seek to identify better parameter settings, algorithm selectors, and portfolios of algorithms, respectively, based on historical data. Although all these methods are within the scope of EBO and thus relevant to one another, they were developed through independent paths and have never been discussed in a unified context. We first bring together the existing literature on the offline scenario of EBO, and review them under the unified umbrella of EBO, so as to make the key issues for EBO clearer.

(2) A new offline training approach for EBO. A (and probably the most fundamental) form of EBO is to train it with many instances of the problem of interests, so as to obtain a well-developed solver before deployment. This scenario involves at least two questions, i.e., where comes the training instances and how the solver is adapted (trained) to the training instances. These two issues were usually treated through two independent phases and seldom addressed simultaneously in the literature. We argue that they are inter-correlated and propose a co-evolutionary framework, namely LiangYi to address them as a whole.

(3) A case study of LiangYi on Travelling Salesman Problem (TSP). To assess the potential of LiangYi, a specific instantiation of it is implemented based on the Chained Lin-Kernighan (CLK) algorithm for the TSP. Empirical studies are conducted to compare LiangYi to other state-of-the-art methods for fine-tuning CLK algorithm, as well as to investigate the properties of LiangYi.

The rest of this paper is structured as follows. Section 2 first gives a definition

of the offline training in EBO and presents the key issues for describing the offline training methods, and then review the existing methods. Section 3 presents the approach LiangYi. Section 4 instantiates LiangYi on the TSP and reports the empirical results. Section 5 concludes the paper and outlines directions for future research.

2. Offline Training Methods in EBO

Given an optimization task T and a performance metric m, the training in EBO is defined as making a solver s improve at optimization task T with respect to performance metric m through experience E. This definition borrows some basic concepts (i.e., T, m, E) from the definition of machine learning by Mitchell [**?**], yet each of them has a concrete meaning here. Specifically, the optimization task T is conceptually an instance set containing all the target instances to which the output solver is expected to be applied. The performance metric m is user-specified and it is often related to the computational resources consumed by the solver (such as runtime or memory) or the quality of the solution found. Basically, the experience E that is produced in the training is the processes of the solver s solving training instances.

To improve the solver s at task T, a training method must consider two things: How to use the solver s and the training instances to produce useful experience and how to exploit the experience so as to enhance the solver s. Generally, a solver s is comprised of multiple different parts (for example, a search-based solver includes at least an initialization module and a search operator). It is conceivable that if the solver s is improved by the training, some parts of s are necessarily changed during the training process. Based on these analyses, we consider that an offline training method in EBO consists of three essential parts:

- the form of the solver being trained;
- the settings of the training instances;

- the training algorithm that manipulates the solver and the training instances to produce experience, and exploits the experience to improve the solver.

With this framework, we can describe offline training methods in EBO in a unified way. In the combinatorial optimization field, there have been various attempts by different communities to obtain solvers through training ways. The next few sections review these researches with the presented framework in turn.

2.1. Automatic Algorithm Configuration Methods

The first class of methods are automatic algorithm configuration (AAC) methods [**?** **?**]. AAC methods improve the solver (a parameterized algorithm) at the optimization task by finding parameter values under which the solver achieves high performance on the target instances. Specifically, AAC methods adopt a two-stage strategy. They first build a training set containing the training instances that are representative of the target instances, and then run the training algorithms to find high-performance parameter values on the training set. Due to the similarity between the training instances and the target instances, the found parameter configurations are expected to perform well on the target instances as well. A number of efficient methods have been developed in the field of AAC, such as CALIBRA [**?**], ParamILS [**?**], GGA [**?**], SMAC [**?**] and irace [**?**]. With the framework presented previously, AAC methods are expressed as follows:

- The solver s being trained is a parameterized algorithm. [1]
- The efficacy of AAC methods depends greatly on the selection of the training instances, that is, the training instances should represent the target

[1] Although there may be some significant differences between the parameterized algorithms (in the aspects such as the types of the parameters, or the number of the parameters) that different AAC methods can handle, we choose to ignore these details because what we want to clarify here is which part of the solver is changed by the training, and the solver description, i.e., a parameterized algorithm, is enough for this purpose. Such a simplicity principle also applies in the reviews of other kinds of methods.

instances well so that the optimized performance on the training instances could be favourably transferred to the target instances. The usual practice in setting training instances for *AAC* methods [**?** **?** **?** **?**] is that the training instances are directly selected from some benchmarks, or are randomly generated through some instance generators. Such practice is based on the assumption that the selected benchmarks and generators could represent the target scenarios to which the output solver will be applied to, particularly some practical application scenarios. This assumption however has sparked some controversy [**?** **?**], which we will discuss more in Section 3.

- Essentially, in the training process, *AAC* methods run training algorithms that keep testing different parameter configurations with the training set and obtaining the testing results. The experience E produced in the training process is actually those tested parameter configurations and the corresponding testing results. The way of exploiting E so as to enhance the solver is straightforward — using the best-performing one as the parameter configuration of the output solver (so the parameter configuration of the solver is the part that is changed by the training).

  When producing the experience, the training algorithms mainly need to address such a problem: Which parameter configurations should be evaluated? Various *AAC* methods have proposed different solutions to this problem. Among them the most notable ones include using the experimental design methods to generate candidate parameter configurations, such as the case in CALIBRA [**?**], searching the parameter space with a meta-heuristic, such as the case in ParamILS [**?**] (which uses an iterated local search) and GGA [**?**] (which uses a genetic algorithm), and adopting a model to bias the generation of candidate parameter configurations, such as the case in SMAC [**?**] (which builds an explicit regression model to describe the dependence of algorithm performance on parameter settings) and irace [**?**] (which builds a probability model over the parameter space). In addition to determining the candidate parameter configurations

to be evaluated, some other important issues need to be addressed include which instances are used to evaluate a parameter configuration and when to terminate the evaluation of those poorly performing configurations. A detailed review of these aspects in this area is beyond the scope of this paper and one may refer to [**?** **?**] for a more comprehensive treatment on the subject.

2.2. Portfolio-based Algorithm Selection Methods

The second class of methods are portfolio-based algorithm selection (PAS) methods [**?** **?** **?**]. Although there are different interpretations of this term "portfolio" in the literature, we use it here to denote a solver that contains several candidate algorithms and always selects one of them when solving a problem instance. To improve the solver (an algorithm portfolio), unlike AAC methods, PAS methods do not change the algorithms that constitute the solver, but build a selector that can accurately select the best from the candidate algorithms for each instance. PAS methods adopt the same two-stage strategy as AAC methods, except that the second-stage training algorithms in PAS methods are used to establish the selector. With the framework presented previously, PAS methods are expressed as follows:

- The solver s being trained is an algorithm portfolio.
- The training instance settings for PAS methods are the same as AAS methods.
- Although the training algorithms adopted by different PAS methods are different in detail, their basic ideas are mostly the same. In the training process, the training algorithms first gather performance data of the candidate algorithms on the training instances, and then build an algorithm selector based on the gathered data. The experience E produced in the training is the performance data, which is obtained by running each candidate algorithm on all training instances. The exploitation of E is carried out in this way: First suitable features that characterize problem

instances are identified, and then the feature values of training instances are computed; Once each training instance is represented by a vector of feature values, the performance data is transformed into a set of training data, and based on these training data, machine learning techniques are used to learn a mapping from instance features to algorithms. This mapping is exactly the algorithm selector that will be used in the output portfolio (so the algorithm selector of the solver is the part that is changed by the training).

Various PAS methods build different models to do the mapping (selection), such as regression models [**?** **?** **?** **?** **?** **?**] (so-called empirical performance models), classification models [**?** **?** **?** **?** **?** **?**] and ranking models [**?**]. [2] For additional information one may refer to [**?**] which presents a survey of many approaches for algorithm selection from cross-disciplinary perspectives and [**?**] which presents a survey focusing on the contributions made in the area of combinatorial search problems.

2.3. Automatic Portfolio Construction Methods

The third class of methods are automatic portfolio construction (APC) methods, which seek to automatically build an algorithm portfolio based on a parameterized algorithm from scratch. An APC method can be seen as a combination of an AAC method and a PAS method, that is, APC methods simultaneously change the member algorithms and build the algorithm selectors. There are two representative methods, dubbed Hydra [**?** **?**] and ISAC [**?** **?**], of this category. Another class of portfolio construction methods, called automatic parallel portfolio construction ($APPC$) methods, differ from APC methods in that they seek to construct a *parallel* algorithm portfolio that runs all candidate algorithms in parallel when solving an instance. In other words, $APPC$ methods also change the member algorithms of the solver, which makes them closely

[2] Although many of the methods cited here are not originally proposed for combinatorial optimization problems, the ideas behind them are very general and apply to the combinatorial optimization problems as well.

related to AAC methods, but do not involve any algorithm selection. The representative methods of this category is ParHydra [**?**]. Like AAC methods and PAS methods, both APC methods and $APPC$ methods also adopt a two-stage strategy. With the framework presented previously, APC methods and $APPC$ methods are expressed as follows:

- The solver s being trained by APC methods is an algorithm portfolio, while the solver s being trained by $APPC$ methods is a parallel algorithm portfolio.

- The training instance settings for APC methods and $APPC$ methods are the same as AAC methods.

- During the training process, APC methods call an AAC method and a PAS method as subroutines, while $APPC$ methods only call an AAC method as a subroutine. Thus the experience collection and exploitation in APC methods and $APPC$ methods are done by their respective subroutines.

  Different APC methods mainly differ in how they specifically arrange the calls of the AAC algorithm and the PAS algorithm. Hydra [**?**] iteratively runs them to build the portfolio. In the first iteration, it uses the AAC algorithm to identify a parameter configuration with the best overall performance on the training set, and in each subsequent iteration, it uses the AAC algorithm to find a parameter configuration that maximizes marginal performance contribution across the configurations identified in the previous iterations. These marginal contributions are calculated relative to the current algorithm selector, which is built by the PAS algortihm. A later version of Hydra [**?**] aims to find multiple algorithm configurations at each iteration to speed up the training process. Unlike Hydra, ISAC [**?**] runs the PAS algorithm and the AAC algorithm sequentially. It first uses a clustering algorithm (the PAS algorithm) to divide training instances into clusters based on normalized instance features, thus building an algo-

rithm selector that is actually a 1-Nearest Neighbor model (this means at runtime, for each instance the nearest cluster is determined and the corresponding algorithm is invoked). Then ISAC applies the AAC algorithm to find a good algorithm configuration for each cluster. The later version of ISAC [**?**] extends ISAC by using the AAC algorithm on a set of different parameterized algorithms to find a good algorithm configuration for each instance cluster.

The representative $APPC$ method ParHydra [**?**], like Hydra, also calls the AAC algorithm iteratively. In each iteration ParHydra aims to find an algorithm configuration that maximizes marginal contribution. Unlike Hydra, ParHydra has no PAS algorithm involved and the marginal contribution is calculated relative to the current portfolio running in parallel. A variant of ParHydra, dubbed $\text{ParHydra}_{\text{b}}$, which is also proposed by [**?**], extends ParHydra by using the AAC algorithm on a set of different parameterized algorithms to identify an algorithm configuration in each iteration.

2.4. Transfer Methods

The last class of methods are transfer methods, which explicitly extract useful information from solving processes of the training instances, and use it to improve the performance of the solver on the target instances. The biggest difference between transfer methods and the methods reviewd above is that the formaer collect experience based on every single instance, while the latetr collect based on a set of instances. Recall that AAC methods run the candidate algorithms on a set of instances to evaluate them, and PAS methods run the member algorithms on all training instances to collect the training data. On comparasion, transfer methods extract information from each individual training instance, and the extraction of different instances is independent of each other. Two specific transfer methods, dubbed XSTAGE [**?**] and MEMETrans

[**?**], are reviewed here [3]. Both of them enhance search-based solvers by introducing high-quality solutions during the solving processes. With the framework presented previously, they are expressed as:

- The solver s being trained by XSTAGE is a multi-restart local search algorithm in which a value function is used to determine the starting point for each local search process.

  The solvers s being trained by MEMETrans are search-based meta-heuristics.

- XSTAGE collects training instances and target instances from the same benchmark, implicitly assuming these instances are similar enough to make the transfer useful. MEMETrans collects training instances and target instances from multiple heterogeneous benchmarks, and uses an instance similarity measure to determine given a target instance, which training instances will be used.

- Both XSTAGE and MEMETrans first run the solvers to solve the training instances, producing the experience that is actually the solving processes of the training instances. They differ in the exploitation of the experience. XSTAGE extracts a value functions from the solving process of each training instance, and combine all of them into one value function that determines starting solution for the solver. MEMETrans learns a mapping from each solved training instance to corresponding optimal solution, and then combines all of them to be an initialization module that helps generate high-quality initial solutions for the solver.

3. LiangYi

The two main issues of offline training are selecting the training instances and training the solvers. The two-stage strategy, which is widely adopted by

[3] In the paper [**?**] that proposed XSTAGE, the term XSTAGE is used to denote the composition of the training method and the solver. We use XSTAGE here only to denote the training method. The term MEMETrans is created by us to denote the transfer method proposed by [**?**] since the paper does not give a term for this method.

existing training methods such as *AAC* methods [? ? ? ? ?], *PAS* methods [? ? ? ? ? ? ? ? ? ? ? ? ?], *APC* methods [? ? ? ? ?] and transfer methods [? ?], treat them in two independent phases. However, intrinsically, these two issues are correlated. From the perspective of developing solvers, the greatest chance of the solver getting improved is on those problem instances which the current solver cannot solve well [4]. Thus those hard instances for the solver are best suited as the training instances. On the other hand, during the training process, the solver is being adapted to the training instances; thereby as the training proceeds, instances that are previously appropriate for the solver may not be appropriate any more, in which case fixing the training instances (as two-stage strategy methods do) is actually not helpful in improving the solver, but will result in the waste of computing resources. A better strategy is to dynamically changing the training instances during the training process to keep them always appropriate (hard) for the solver, so that the solver can be continuously improved.

As aforementioned, usually the training instances for two-stage methods are directly selected from some benchmarks [? ?], or are randomly generated through some instance generators [? ? ? ? ? ? ? ? ? ?], based on an assumption that the benchmarks and generators could represent the target scenarios to which the solver is expected to be applied. However, such assumption is not always tenable. On the contrary, in many cases, the commonly studied benchmark instances and randomly generated instances lack diversity, are too simple and rarely resemble real-world instances [? ?]. Such risks could be avoided by dynamically changing the training instances. First, this strategy selects training instances that are never easy for the solver. Second, this strategy keeps changing the training instances, which naturally introduces the diversity.

Based on the above considerations, we propose a new training method,

[4] This is also the key idea behind Hydra [? ?], in which the training instances are given different importance according to their hardness to the solver. However, such adjustment of instance importance is still within a fixed training set.

dubbed LiangYi [5]. Basically, LiangYi is a competitive co-evolutionary [?] framework that alternately trains the solver and searches for new training instances. It maintains a set of algorithms and a set of instances, and each set strives to improve itself against the other during the evolutionary process. The details of this framework are elaborated below.

3.1. Design Principles

3.1.1. Basic Solver

The form of the solver being trained by LiangYi is a parallel portfolio that runs all candidate algorithms in parallel when solving an instance [6]. In the co-evolutionary framework, the parallel portfolio is called an algorithm population. The reason for choosing an algorithm population other than a single algorithm as the solver is simple: An algorithm population has the potential to achieve better overall performance than a single algorithm [? ? ? ?], since it is often the case that for a problem domain there is no single best algorithm for all possible instances, but different algorithms perform well on different problem instances.

3.1.2. General Framework

Overall LiangYi consists of two components — The first component is a training module for training the solver, and the second component is a sampling module which samples the target instance space for updating the training instances. These two modules are coupled in a co-evolutionary framework. Specifically, LiangYi maintains two populations in which one population is an algorithm population (the solver) and the other is an instance population (the

[5] The name "LiangYi" comes from the Taoism of Chinese philosophy. Generally, it means two opposite elements of the world that interact and co-evolve with each other.

[6] Employing parallel solvers to problem instances is an emerging area in training solvers [?]. Note that running all algorithms in parallel is different from an algorithm portfolio [? ? ?], which typically involves some mechanism (e.g., selection [? ?] or scheduling [?]) to allocate computational resource to different algorithms. Running multiple algorithms in parallel does not require any resource allocation as all involved algorithms are assigned with the same amount of resource. This may lead to the waste of resource to some extent but can keep the implementation simple.

training set). The training module and the sampling module are as two sub-procedures that are performed alternatively by LiangYi to train the algorithm population and to adjust the instance population, respectively. From the perspective of evolution, LiangYi is a process in which two populations are evolved alternatively. At each iteration, LiangYi first evolves the algorithm population (using the training module) to improve the performance of the algorithm population on the instance population while keeping the good performances obtained by the algorithm population in previous iterations, and then evolves the instance population (using the sampling module) to discover and include those instances that cannot be solved well by the algorithm population, and then enters the next iteration.

Intuitively, if we consider an instance in the target instance space "covered" by a solver as it can be solved well by the solver, the essence of LiangYi is to enlarge the algorithm population's coverage on the target instance space by a) making the algorithm population cover the area that has not been covered yet and b) keeping the algorithm population covering the area that has already been covered. Figure 1 gives an intuitive visual example. In the example the algorithm population before training contains member algorithms that are neither good nor complementary, and after training, the algorithm population obtain new member algorithms that are both individually better and more complementary, which enlarges the coverage of the algorithm population a lot on the basis of the original one.

3.2. Implementation Details

For the sake of brevity, henceforth we use abbreviations for the frequently used terms. All the used abbreviations and corresponding terms are listed in Table 1.

The pseudo code of LiangYi is outlined in Algorithm Framework 1. LiangYi first randomly initializes AP as AP_1 and IP as IP_1 (Lines 1-2) and then enters the iteration in which AP and IP are evolved alternatively (Lines 4-8). During the k-th iteration, LiangYi first evolves AP from AP_k to AP_{k+1}, then evolves

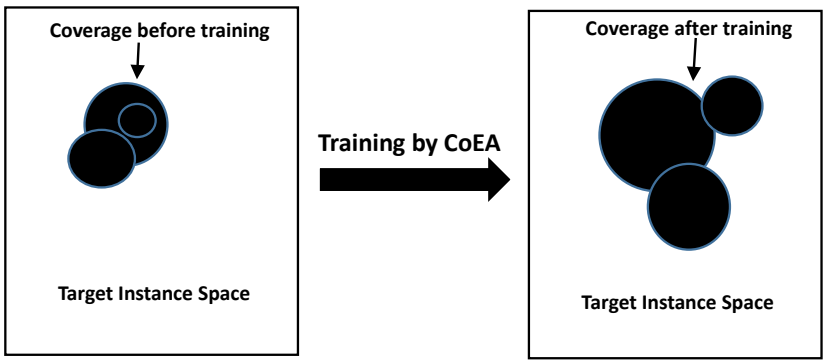

Figure 1: An example of LiangYi enlarging the coverage of the algorithm population on the target instance space. The target instance space is the solid square. The algorithm population contains 3 member algorithms. The coverage of each algorithm on the target insatnce space is indicated by a black circle.

Table 1: Aberrations and corresponding terms

AP	Algorithm population
IP	Instance population
AP_k	Algorithm population at the initial stage of the k-th iteration of LiangYi
IP_k	Algorithm population at the initial stage of the k-th iteration of LiangYi
alg	An algorithm belonging to the algorithm population
ins	An instance belonging to the instance population
$EvolveAlg$	Training module of LiangYi
$EvolveIns$	Sampling module of LiangYi
N_{AP}	Number of algorithms in algorithm population
N_{IP}	Number of instances in instance population
$params_{AP}$	Parameter set that controls training module
$params_{IP}$	Parameter set that controls sampling module
$params_R$	Parameter set that controls $RemoveWorst$ procedure
$P(solver, instanceset)$	Performance of the solver on the instance set
$Aggr()$	Aggregate function
M	Performance matrix of algorithm population on instance population
$f_{AP}(alg)$	Fitness of algorithm alg
$f_{IP}(ins)$	Fitness of instance Ins

IP from IP_k to IP_{k+1}. After the k-th iteration, LiangYi enters the $(k+1)$-th iteration with AP as AP_{k+1} and IP as IP_{k+1}. Finally, when LiangYi is terminated, the AP is returned as the output solver (Line 9).

Algorithm Framework 1: LiangYi($N_{AP}, N_{IP}, params_{AP}, params_{IP}$)

Input: Number of algorithms in AP, N_{AP}, number of instances in IP, N_{IP}, set of the parameters which control $EvolveAlg$, $params_{AP}$, set of the parameters which control $Evolveins$, $params_{IP}$

Output: Algorithm population AP_k

1. $AP_1 \leftarrow$ Randomly generate an algorithm population with N_{AP} algorithms
2. $IP_1 \leftarrow$ Randomly generate an instance population with N_{IP} instances
3. $k \leftarrow 1$
4. **while not** $LiangYiTermination()$ **do**
5. $\quad AP_{k+1} \leftarrow EvolveAlg(AP_k, IP_k, k, params_{AP})$
6. $\quad IP_{k+1} \leftarrow EvolveIns(AP_{k+1}, IP_k, params_{IP})$
7. $\quad k \leftarrow k+1$
8. **end while**
9. **return** AP_k

The training module and the sampling module in LiangYi are implemented as two evolutionary procedures $EvolveAlg$ and $EvolveIns$, respectively. In general, any search method can be used in these two procedure. In this work, evolutionary algorithms (EAs) are employed as the off-the-shelf tools, because EAs are directly suitable for handling populations (of either algorithms or instances) and are less restricted by the properties of objective functions [**?**] (in comparison to other search methods such as gradient descent that requires differentiable objective functions). Note that, in the process of LiangYi, when training algorithm population or adjusting the instance population, the objective functions are a priori unknown and are changing in every iteration, so it is a must that the two modules are able to cope with whatever objective landscape encountered. The behaviors of $EvolveAlg$ and $EvolveIns$ are controlled by parameter sets $params_{AP}$ and $params_{IP}$, respectively.

3.2.1. Preliminaries

Now we address three issues before going into the details of $EvolveAlg$ and $EvolveIns$. First of all, we explain how algorithms and instances are repre-

sented. Generally, when applying EAs, the choice of individual representation would determine which genetic operators are applicable. In this paper. a fixed-size linear integer representation is chosen for both algorithms and instances, namely, each algorithm and each instance is represented by an array of integers. Although such representation might restrict the application scope of LiangYi, it makes simple genetic operators [**?**] (e.g., uniform crossover and uniform mutation) directly applicable without any customization. For specific types of algorithms and problems, more complex representations (e.g., tree-based representation) and corresponding genetic operators are needed, which will be a subject of future research.

The second issue that needs to be addressed is to define the algorithms in AP. Similar with automatic portfolio construction methods (see Section 2.3) , LiangYi builds solver (the algorithm population) based on a parameterized algorithm. In other words, each algorithm in AP is actually a parameter configuration which is a sample of the configuration space consisting of all possible configurations of the parameterized algorithm.

The last issue that needs to be addressed is how to measure performance of AP. For simplicity henceforth we use $P(\text{solver}, \text{instance set})$ to denote the performance of a solver on the instance set according to a given performance metric m, in which the solver could be a single algorithm or an algorithm population and the instance set could be a single instance or an instance population. Since a population-based solver runs all member algorithms in parallel when solving an instance, the performance of AP on an instance ins is the best performance achieved by its member algorithms on ins (we assume a larger value is better for m without loss of generality), i.e,

$$P(AP, ins) = \max_{alg \in AP} P(alg, ins). \tag{1}$$

The performance of AP on IP is an aggregated value of the performance of AP

on each instance of IP, i.e.,

$$P(AP, IP) = \underset{ins \in IP}{Aggr} \left(P(AP, ins) \right) \tag{2}$$

where $Aggr()$ is an aggregate function. The performance metric m and the aggregate function $Aggr()$ are user-specified.

3.2.2. Evolution of the Algorithm Population

The pseudo code of $EvolveAlg$ is shown in Procedure 2. First of all, AP_k is

Procedure 2: $EvolveAlg(AP_k, IP_k, k, params_{AP})$

Memory is a global cache which stores the final AP and final performance matrix M of $EvolveAlg$ of each iteration of LiangYi

Input: Algorithm population AP_k, instance population IP_k, number of current iteration of LiangYi, k, parameter set $params_{AP}$ containing: crossover probability cro_{alg}, mutation rate mu_{alg}, set of parameters which control $RemoveWorst$, $params_R$

Output: Algorithm population AP_{k+1}

1. $M \leftarrow$ Test AP_k on IP_k
2. **while not** $EvolveAlgTermination()$ **do**
3. $a_1, a_2 \leftarrow$ Randomly select two individuals from AP_k
4. $a_{new} \leftarrow$ Use a_1, a_2 to generate an offspring by uniform crossover with probability cro_{alg} and uniform mutation with rate mu_{alg}
5. $M_{new} \leftarrow$ Test a_{new} on IP_k
6. $AP'_k \leftarrow AP_k \cup \{a_{new}\}$
7. $M' \leftarrow$ Concatenate M and M_{new}
8. $AP_k, M \leftarrow RemoveWorst(AP'_k, M', k, params_R)$
9. **end while**
10. Add AP_k and M to $Memory$
11. $AP_{k+1} \leftarrow AP_k$
12. **return** AP_{k+1}

tested on IP_k and the result is represented by a $N_{AP} \times N_{IP}$ matrix M (N_{AP} and N_{IP} are the number of the algorithms in AP and the number of the instances in IP, respectively), in which each row corresponds to an algorithm in AP_k and each column corresponds to an instance in IP_k, so peach entry in M is the performance of the corresponding algorithm in AP_k on the corresponding instance in IP_k (Line 1). Then at each generation, two individuals are randomly selected from AP_k and an offspring is generated by using uniform crossover and

mutation [**?**] (Lines 3-4). The generated algorithm a_{new} is then tested on IP_k and the result is represented by a $1 \times N_{IP}$ matrix M_{new} (Line 5). The last step is to decide the survivors of this generation (Lines 6-8). Together with all the algorithms in AP_k, a_{new} is put into a temporary algorithm population AP'. The corresponding performance matrix M' to AP', of which the size is $(N_{AP} + 1) \times N_{IP}$, is constructed by concatenating M and M_{new}. Then the procedure $RemoveWorst$ is run to decide which algorithm in AP' will be removed.

Basically, $RemoveWorst$ firstly calculates the fitness of each algorithm in AP' and then selects the algorithm with the lowest fitness to be removed. The core of $RemoveWorst$ is its fitness evaluation. The idea is that an algorithm will be preferred only if it contributes to the algorithm population, and the more it contributes, the more it is preferred. The contribution of a member algorithm is actually the performance improvement it brings to the population, which can be calculated as the population's performance loss caused by removing the algorithm. Formally, let $C(AP, IP, alg)$ denote the contribution of algorithm alg to the performance of AP on IP. If $|AP| > 1$, which means AP contains other algorithms besides alg, $C(AP, IP, alg)$ is calculated via Equation (3):

$$C(AP, IP, alg) = |P(AP, IP) - P(AP - \{alg\}, IP)|, \tag{3}$$

where $P(AP, IP)$ and $P(AP - \{alg\}, IP)$ are calculated via Equation (2). If $|AP| = 1$, which means AP only contains one algorithm (i.e., alg), then removing alg from AP will cause complete performance loss on IP. In this case, $C(AP, IP, alg)$ is calculated via Equation (4):

$$C(AP, IP, alg) = \alpha |P(AP, IP)|, \tag{4}$$

where the parameter $\alpha > 0$.

Based on Equation (4), for each member algorithm alg of the temporary algorithm population AP'_k, the performance contribution on instance popula-

tion IP_k is $C(AP'_k, IP_k, alg)$. A high contribution indicates that alg should be reserved to next generation. However, directly using $C(AP'_k, IP_k, alg)$ as the fitness of alg is not appropriate. As aforementioned (see Section 3.1), the evolution of AP_k should not only improve the performance of AP_k on IP_k, but also keep the good performances obtained in previous iterations (on $IP_1, ..., IP_{k-1}$). Using $C(AP'_k, IP_k, alg)$ to evaluate alg only considers the first target. Hence the fitness of alg is calculated based on two types of contributions: The first one is the current performance contribution, i.e., $C(AP'_k, IP_k, alg)$, and the second one is the historical performance contribution of alg on $IP_1, ..., IP_{k-1}$ (if there are).

To calculate the historical performance contribution of a member algorithm, the concept age is introduced to describe how long the algorithm has been in AP. Suppose a member algorithm alg was added to AP in the j-th iteration of LiangYi (and now is in the k-th iteration), the age of alg is $k - j$. The performances of alg on IP_j,...,IP_k are known because alg has been tested on them in corresponding iterations [7]. To calculate the historical performance contribution of alg on $IP_r (j \leq r \leq k-1)$, the algorithms that satisfy the condition $age \geq (k - r)$ are selected from the temporary algorithm population AP'_k (our target algorithm alg is also selected, since its age is $k - j$, which satisfies the condition) to form a virtual algorithm population $virtualAP_r$. The condition $age \geq (k-r)$ indicates that these selected algorithms were added to AP during or before the r-th iteration of LiangYi, so they have been tested on IP_r. The performances of these algorithms on IP_r are represented by a $|virtualAP_r| \times N_{IP}$ matrix $virtualM_r$. If $|virtualAP_r| > 1$, which means $virtualAP_r$ contains other algorithms besides alg, the performance contribution of alg on IP_r is calculated

[7] Note that the performance contributions of alg on IP_1,...,IP_{j-1} are not considered in this paper because the performances of alg on them are unknown. To obtain these performances, we can store IP_1,...,IP_{j-1} and test alg on them. However, this would make the computational cost and the storage cost increase fast over time.

via Equation (3):

$$C(virtualAP_r, IP_r, alg) = |P(virtualAP_r, IP_r) - P(virtualAP_r - \{alg\}, IP_r)|.$$

If $|virtualAP_r| = 1$, which means $virtualAP_r$ only contains alg, then removing alg from $virtualAP_r$ will cause complete loss of performance on IP_r. In this case, $C(virtualAP_r, IP_r, alg)$ is calculated via Equation (4):

$$C(virtualAP_r, IP_r, alg) = \alpha|P(virtualAP_r, IP_r)|. \quad (4)$$

Now we have all the performance contributions of alg on IP_j,...,IP_k. The fitness of alg, denoted as $f_{AP}(alg)$, is calculated via Equation (5)

$$f_{AP}(alg) = \frac{\beta\Sigma_{j\leq r\leq k-1}C(virtualAP_r, IP_r, alg) + C(AP'_k, IP_k, alg)}{k-j+1} \quad (5)$$

where k is the index of the current iteration of $pLiangYi$, j is the age of alg, and β is a nonnegative parameter. The terms $C(virtualAP_r, IP_r, alg)(j \leq r \leq k-1)$ are historical performance contributions on IP_j,...,IP_{k-1}, while $C(AP'_k, IP_k, alg)$ is the current performance contribution on IP_k. Thereby the numerator in the fraction is actually a weighted sum of $(k-j+1)$ performance contributions, in which the parameter β is used to balance between historical performance contributions(on $IP_1, ..., IP_{k-1}$) and current performance contribution (on IP_k).

The pseudo code of $RemoveWorst$ is demonstrated in Procedure 3. First the fitness of each member algorithm in AP'_k is calculated (Lines 1-15). Specifically, for an algorithm alg which was added to AP in the j-th iteration of LiangYi, $k-j$ virtual algorithm populations, i.e., $virtualAP_j, ..., virtualAP_{k-1}$, are constructed (according to the global cache Memory) to calculate its historical performance contributions on $IP_j, ..., IP_{k-1}$, via Equation (3) or Equation (4) (Lines 2-12). Together with the current performance contribution calculated via Equation (3) (Line 13), the historical contributions are used to calculate the fitness of alg via Equation (5) (Line 14). After the fitness of each algorithm in

AP'_k has been calculated, the algorithm with the lowest fitness will be removed (Line 16).

Procedure 3: $RemoveWorst(AP'_k, M', k, params_R)$

Input: Temporary algorithm population AP'_k, temporary performance matrix M', number of current iteration of LiangYi, k, parameter set $params_R$ containing α (used in Equation (4)) and β (used in Equation (5))
Output: Algorithm population AP_k, performance matrix M

1. **for each** algorithm alg in AP'_k **do**
2. $age \leftarrow$ Query $Memory$ for how many iterations alg has been staying
3. $j \leftarrow k - age$
4. **for** $r \leftarrow j$ to $k-1$ **do**
5. $VirtualAP_r \leftarrow$ Select algorithms which satisfy the condition $age \geq (k-r)$ (according to $Memory$) from AP'_k to form a virtual algorithm population
6. $M_r \leftarrow$ Construct corresponding performance matrix to $VirtualAP_r$ according to $Memory$
7. **if** $|virtualAP_r| > 1$ **then**
8. $C(virtualAP_r, IP_r, alg) \leftarrow$ Calculate the algorithm contribution of alg on IP_r via Equation (3)
9. **else**
10. $C(virtualAP_r, IP_r, alg) \leftarrow$ Calculate the algorithm contribution of alg on IP_r via Equation (4)
11. **end if**
12. **end for**
13. $C(AP'_k, IP_k, alg) \leftarrow$ Calculate the algorithm contribution of alg on IP_k via Equation (3)
14. $f_{AP}(alg) \leftarrow$ Calculate the fitness of alg via Equation (5)
15. **end for**
16. $AP_k \leftarrow$ Remove the algorithm with the lowest fitness from AP'_k
17. $M \leftarrow$ Remove the corresponding row in M' to the removed algorithm
18. **return** AP_k, M

3.2.3. Evolution of the Instance Population

As aforementioned (see Section 3.1) the evolution of IP aims at discovering those instances that cannot be solved well by AP; thus the fitness of an instance in IP is measured by how AP performs on it — the worse the performance, the higher the fitness.

The pseudo code of $EvolveIns$ is demonstrated in Procedure 4. First of all, AP_{k+1} is tested on the IP_k and the result is represented by a $N_{AP} \times N_{IP}$ matrix

M (Line 1), and the fitness of each instance is calculated (Lines 2-4). The fitness of an instance ins, denoted as $f_{IP}(ins)$, is calculated via Equation (6):

$$f_{IP}(ins) = -P(AP_{k+1}, ins), \tag{6}$$

where $P(AP_{k+1}, ins)$ is the performance of AP_{k+1} on instance ins, calculated via Equation (1). At each generation, $N_{IP} * re$ new instances are generated by repeatedly selecting two instances from IP_k and creating two offsprings using uniform crossover with probability cro_{ins} and uniform mutation with rate mu_{ins} [**?**](Lines 6-11). These offsprings are then tested against the algorithm population AP_{k+1} and the fitness of each offspring is calculated (Lines 12-15). At the end of this generation, all instances in IP_k and the offsprings are put into a candidate pool and the worst $N_{IP} * re$ instances are removed (Lines 16-19).

4. Case Study: the Travelling Salesman Problem

The main purpose of this section is to empirically verify whether LiangYi is an effective method for training solvers. We evaluated LiangYi on the Travelling Salesman Problem (TSP) [**?**], one of the most well-known computationally hard optimization problem. Specifically, the symmetric TSP, i.e., the distance between two cities is the same in each opposite direction, with Euclidean distances in a two-dimensional space is considered here. In the remainder of this section, we first give the target scenario (including the optimization task and the performance metric) where LiangYi is applied, and then instantiates LiangYi for the scenario. After that, we first compare LiangYi with other existing training methods, and then we investigate the properties of LiangYi to see whether it is able to perform as expected in design.

4.1. Target Scenario

The optimization task T considered here are all TSP instances with problem size equal to 500, i.e., the number of cities equals to 500. The number of the

Procedure 4: $EvolveIns(AP_{k+1}, IP_k, params_{IP})$

Input: Algorithm population AP_{k+1}, instance population IP_k, parameter set $params_{IP}$ containing crossover probability mu_{alg} , mutation rate mu_{ins} and replacement ratio res

Output: Instance population IP_{k+1}

1. $M \leftarrow$ Test AP_{k+1} on IP_k
2. **for each** instance ins in IP_k **do**
3. $\quad f_{IP}(ins) \leftarrow$ Calculate the fitness of ins via Equation (6)
4. **end for**
5. **while not** $EvolveInsTermination()$ **do**
6. $\quad offsprings \leftarrow \varnothing$
7. $\quad$ **for** $i \leftarrow 1$ to $\frac{N_{IP} * res}{2}$ **do**
8. $\quad\quad ins_1, ins_2 \leftarrow$ Select two parents from IP_k with tournament selection
9. $\quad\quad ins_{new1}, ins_{new2} \leftarrow$ Use ins_1, ins_2 to generate two offsprings by using uniform crossover with probability cro_{ins} and uniform mutation with rate mu_{ins}
10. $\quad\quad offsprings \leftarrow offsprings \cup \{ins_{new1}, ins_{new2}\}$
11. $\quad$ **end for**
12. $\quad M_{offsprings} \leftarrow$ Test AP_{k+1} on $offsprings$
13. $\quad$ **for each** instance ins in $offsprings$ **do**
14. $\quad\quad f_{IP}(ins) \leftarrow$ Calculate the fitness of ins via Equation (6)
15. $\quad$ **end for**
16. $\quad candidates \leftarrow IP_k \cup offsprings$
17. $\quad M_{candidates} \leftarrow$ Concatenate M and $M_{offspring}$
18. $\quad IP_k \leftarrow$ Remove the worst $N_{IP} * re$ instances from $candidates$
19. $\quad M \leftarrow$ Remove the corresponding columns to those removed instances from $M_{candidates}$
20. **end while**
21. $IP_{k+1} \leftarrow IP_k$
22. **return** IP_{k+1}

target instances are theoretically infinite. Training solvers on such kind of tasks have not been studied before.

This work focuses on optimizing the applicability on the target instances, i.e., the performance metric m is applicability. A solver is said to be applicable to an instance if it can find a good enough solution to this instance within a given time. For TSP, the goodness of a solution sol is measured by the percentage by which the tour length of sol exceeds the tour length of the optimum $sol^\star$ [8], abbreviated as PEO(percentage excess optimum):

$$PEO = \frac{lengh(sol) - lengh(sol^\star)}{lengh(sol^\star)} * 100\%.$$

With the definition of PEO, given a cut-off time t, a solver is said to be *applicable* to an instance ins if the PEO of the best solution found by the solver in time t is below a threshold θ. With the definition of the applicability of a solver to a single instance, the applicability of a solver to an instance set is defined as the proportion of the instances to which the solver is applicable.

In this paper very radical values for the cut-off time t and the PEO threshold θ are adopted ($t = 0.1s$ and $\theta = 0.05\%$) to see whether LiangYi is able to train solvers that can work well under such harsh conditions.

4.2. Instantiating LiangYi

In order to instantiate LiangYi for the above scenario, there are several issues to be addressed. The first issue is to specify the performance function $P(\text{solver}, \text{instance set})$ used by LiangYi (see Section 3.2.1) so that LiangYi can optimize the applicability appropriately. The performance of an algorithm alg on an instance ins, i.e., $P(alg, ins)$ in Equation (1), is specified as follow:

$$P(alg, ins) = \begin{cases} 1, & \text{if } alg \text{ is applicable to } ins \\ 0, & otherwise. \end{cases}$$

[8] The optimum $sol^\star$ is obtained using Concorde [**?**], a branch-and-cut based exact TSP solver.

Table 2: The parameters of the CLK used

Parameters	Parameter Type	# of Candidate Values
Initialization Strategy	Categorical	4
Perturbation Strategy	Categorical	4
Search Depth	Numerical	6
Search Width	Numerical	8
Backtrack Strategy	Categorical	14

Intuitively, an algorithm population AP is said to be applicable to an instance ins if any member algorithm of AP is applicable to ins. With $P(alg, ins)$ specified as above, this definition is equivalent to the definition given by Equation (1), namely, AP is applicable to ins if the best member algorithm of AP is applicable to ins. The aggregate function $aggr()$ in Equation (2) is specified as returning mean value of the aggregated terms:

$$P(AP, IP) = \frac{\sum_{ins \in IP} P(AP, ins)}{|IP|},$$

which essentially calculates the proportion of the instances to which AP is applicable.

The second issue is to choose a parameterized algorithm for LiangYi to build an algorithm population based on it. The choice of the parameterized algorithm in this work is Chained Lin-Kernighan (CLK) [**?**]. It is a variant of the Lin-Kernighan heuristic [**?**], one of the best heuristics for solving symmetric TSP. CLK chains multiple runs of the Lin-Kernighan algorithm to introduce more robustness in the resulting tour. Each run starts with a perturbed version of the final tour of the previous run. We extended the original implementation of CLK to allow a more comprehensive control on its components. The parameters of the resulting algorithm are summarized in Table 2. To handle the randomness of CLK, we adopt a simple way - fixing the random seed of CLK and turning it into a deterministic algorithm.

The third issue is to concretely represent the algorithms and the instances

Table 3: The parameter settings of the instantiation of LiangYi for TSP

$EvolveAlg$	$EvolveIns$
$cro_{alg} = 0.6$	$cro_{ins} = 1$
$mu_{alg} = 0.6$	$mu_{ins} = 0.8$
$\alpha = 2$	$res = 0.3$
$\beta = 2$	

with fixed-size integer arrays (see Section 3.2.1). Each algorithm in AP is represented by a list containing 5 integers, each of which indicates its value for the corresponding parameter. Each instance in IP is represented by a list of 500 (x, y) coordinates on a $10^6 \times 10^6$ grids. The random initialization, uniform crossover, and mutation are implemented as conventional ways for fixed-size-list based representations [**?**]. Specifically, the random initialization works by uniform randomly selecting a value (i.e., two coordinates for the instance, a parameter value for the algorithm) from candidate values for each entry of the individual (the algorithm or the instance). The uniform crossover operates by choosing for each entry of the offspring with equal probability either the value of the entry from the first or the second parent. The mutation consists of replacing the value of each entry of the offspring, with a probability, with uniform randomly chosen one from the candidate values.

The last issue is to set the termination conditions and the parameters of LiangYi. The termination condition for LiangYi is the number of iterations reaching 3. In each iteration, procedure $EvolveAlg$ will be run for 500 generations and $EvolveIns$ will be run for 10 generations. The number of algorithms in AP, i.e., N_{AP}, is set to 6, and the number of instances in IP, i.e., N_{IP}, is set to 150. The parameter settings of LiangYi are listed in Table 3. In order to build an algorithm population with complementary algorithms, the diversity between these member algorithms need to be kept on a high level. For this reason, the mutation rate in *evolvealg* is set to a high value (0.4). For *evolveins*, it is important to keep the instance population exploring the target instance space instead of stagnating in some local areas, and therefore the mutation rate

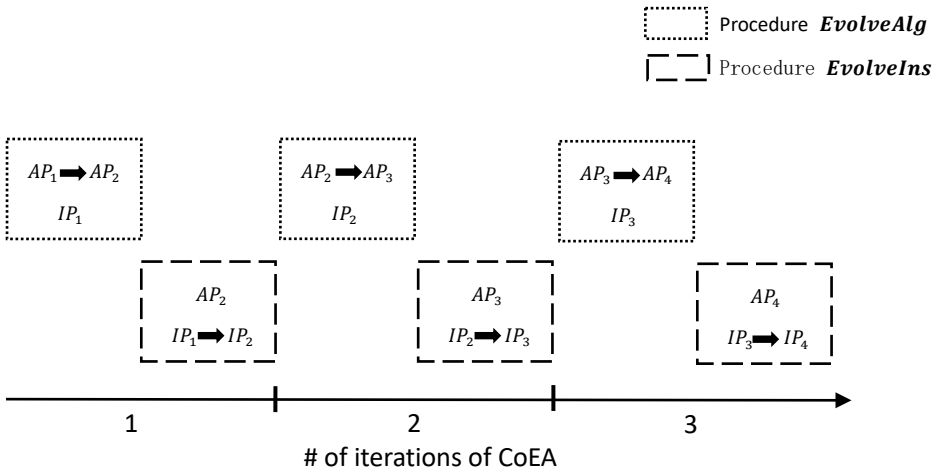

Figure 2: The training process of LiangYi in which AP and IP are evolved alternatively. Dotted rectangle and dashed rectangle represent procedure $EvolveAlg$ and procedure $EvolveIns$, respectively.

in *evolveins* is also set to a high value (0.7).

We applied the instantiation of LiangYi described above to the considered scenario. The training process of LiangYi in which AP and IP are evolved alternatively is depicted in Figure 2.

4.3. Comparative Study

In this section we compare LiangYi with other existing training methods in the considered scenario. Since the algorithm population built by LiangYi is actually a parallel portfolio, we chose ParHydra [**?**] (see Section 2.3) , the state-of-the-art automatic parallel portfolio construction method, as the method to compare with.

4.3.1. Settings of ParHydra

ParHydra accepts a parameterized algorithm to be configured, a set of training instances, and a performance metric to be optimized. For the target scenario considered, the performance metric is applicability. The parameterized algorithm fed to ParHydra is CLK, same as LiangYi. We used two different ways to construct the training sets for ParHydra. The first training set IP_{random}

was built according to the usual practice for two-stage methods — randomly generating a set of instances. Specifically, each instance in IP_{random} was generated by randomly choosing two coordinates for each city on a $10^6 \times 10^6$ grids. The second training set $IP_{training}$ was built by collecting the instance populations which were produced by LiangYi and once served as the training instances during the training process, i.e., $IP_{training} = IP_1 \cup IP_2 \cup IP_3$. Since LiangYi produces instance populations as by-products, it is interesting to see, how good these instance populations are as training instances for existing methods like ParHydra. Both sets contain 450 instances.

ParHydra is an iterative method, which builds the portfolio from scratch and adds an algorithm to it in each iteration (See Section 2.3). Thus we set the iteration number of ParHydra as 6 to keep in line with LiangYi in terms of algorithm number. At each iteration of ParHydra, 15 copies of ParamILS [**?**] were run with different random seeds in parallel to identify an algorithm that maximizes marginal performance contribution across the algorithms identified in the previous iterations. The solvers output by ParHydra based on IP_{random} and $IP_{training}$, are denoted as PH_{random} and $PH_{LiangYi}$ respectively.

4.3.2. Experimental Protocol

Since LiangYi and ParHydra are both stochastic methods, we ran each comparative method 20 times. Specifically, first we ran LiangYi 20 times. Then we obtained 20 $IP_{training}$ from the 20 runs of LiangYi and randomly generated 20 different IP_{random}, and then based on each of these 40 training set, we ran ParHydra to obtain a solver. Finally we obtained 20 AP_4 (the algorithm population output by LiangYi), 20 PH_{random} and 20 $PH_{LiangYi}$.

In order to adequately test the output solvers on target instance space, we generated a huge testing set, denoted as $IP_{testing}$, which contains 10000 TSP instances with the number of cities equal to 500. Specifically, each instance was generated by randomly choosing two coordinates for each city from the interval $[0, 10^6)$. To our knowledge, this is the first time that a testing set of such a large size (10000) is used to test TSP solvers.

All of our experiments were carried out on a workstation of a Xeon CPU with 12 cores and 24 threads at 2.50GHz, running Ubuntu Linux 16.04. The runtime requirements in CPU days were as follows: 120 days for CoEA training (including 20 runs); 600 days for ParHydra training (including 20 runs); 4 days for testing.

4.3.3. Experiment Results

The averaged testing results of each type of solvers are presented in Table 4. We performed 2-sample t-test to compare the results each type of solvers, considering p-values below 0.05 to be statistically significant. The ranking in Table 4 is AP_4, $PH_{LiangYi}$, PH_{random}. $PH_{LiangYi}$ obtained better results than PH_{random}, which indicates the training instances produced by LiangYi are more diverse than random generated ones. It is a little surprising to see AP_4 obtained better results than $PH_{LiangYi}$ at first sight. Different from AP_4, $PH_{LiangYi}$ was trained based on the whole $IP_{training}$, which was produced by LiangYi iteration by iteration; thus it is expected that $PH_{LiangYi}$ would obtain better results than AP_4 on $IP_{training}$ (actually, their performances on $IP_{training}$, i.e., $P(AP_4, IP_{training})$ and $P(PH_{LiangYi}, IP_{training})$, are 0.6063 and 0.6644). The reason why AP_4 obtained better results than $PH_{LiangYi}$ on $IP_{testing}$ is as follow: The dynamic instance sampling strategy used by LiangYi is actually a filter that only keeps those hard instances for AP to make the training focus on them , which makes the actual coverage of AP on the target instances far greater than its coverage on $IP_{training}$ ($0.7001 > 0.6063$) because those easy target instances which are sampled by $EvolveIns$ and are actually covered by AP are all filtered out. Compared to LiangYi, ParHydra accepts all the training instances and only focuses on the training set. The lack of the instance dynamic sampling makes the performance of the output solver greatly depend on how much the training set can represent the target instance space.

Table 4: Averaged performances of $AP_4, PH_{random}, PH_{LiangYi}$ on $IP_{testing}$. All results are presented in terms of applicability, i.e., the proportion of the instances to which the solver is applicable. We performed 2-sample t-test to compare their results, considering p-values below 0.05 to be statistically significant. The second row provide the results of the test, where $'W - D - L'$ indicates how many competitors the corresponding solver is superior to, not significantly different from or inferior to.

AP_4	PH_{random}	$PH_{LiangYi}$
0.7001	0.6600	0.6719
2-0-0	0-0-2	1-1-0

4.4. Investigating the Properties of LiangYi

As aforementioned, the idea behind LiangYi is to optimize the performance of AP on target instances by a) improving its performances on those instances on which it performs badly and b) keeping its good performances on those instances on which it performs well. The main purpose of this section is to investigate whether LiangYi is able to accomplish the two objectives listed above. Specifically, the verification is divided into two parts — the training part and the testing part. In the training part we investigate that, in the training process, whether LiangYi gives satisfactory answers to the following three questions:

(1) Whether is procedure *evolvealg* able to improve the performance of AP on current IP?

(2) Whether is procedure *evolveins* able to degrade the performance of AP on current IP?

(3) Whether is procedure *evolvealg* able to keep the performance of AP on previous IPs?

The second question indicates whether the evolution of IP is able to discover and include hard-to-solve instances to AP, and the first question indicates whether the evolution of AP is able to improve the performance of AP on the hard instances included in the current IP. The combination of these two checks the whether LiangYi is able to accomplish the first objective. The third question checks whether LiangYi is able to accomplish the second objective. In addition to focusing on the three specific aspects, we also directly check if LiangYi is

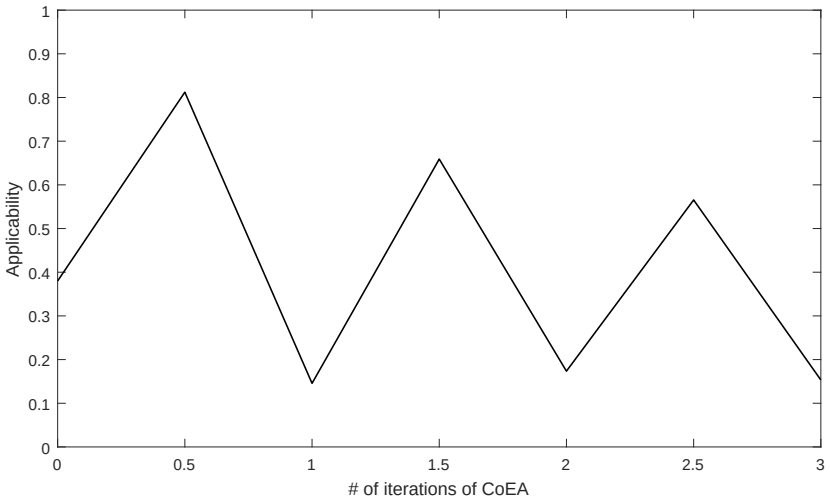

Figure 3: The performance of AP on IP at the beginning ($EvolveAlg$ begins), the middle ($EvolveAlg$ ends) and the end ($EvolveIns$ ends) of each iteration of LiangYi. The performances are represented in terms of applicability, i.e., the proportion of the instances to which AP is applicable, and are averaged over 20 runs.

able to continuously improve AP in the training part. Specifically, we check whether the performances of AP on $IP_{training}$ that are produced in the training are improved by LiangYi. Similarly, in the testing part we also directly check whether the performance of AP at the optimization task is being improved by LiangYi.

4.4.1. Training Part

To answer the first question and the second question, the performances of AP on IP during the training process averaged over 20 runs are plotted in Figure 3. The results depicted in Figure 3 clearly show that, at each iteration of LiangYi, $EvolveAlg$ improves the performance of AP_k on IP_k, and $evolveins$ degrades the performance of AP_{k+1} on IP_k, which gives positive answers to the first two questions, thus confirming the first aspect of the idea behind LiangYi.

The third question is answered in this way: Since procedure $evolvealg$ evolved AP_k to AP_{k+1} to improve the performance on IP_k, we checked whether the improvement from $P(AP_k, IP_k)$ to $P(AP_{k+1}, IP_k)$, i.e., $|P(AP_k, IP_k) - P(AP_{k+1}, IP_k)|$, was kept in subsequent iterations of LiangYi. Specifically,

Table 5: Performances of AP_k,...AP_4 on IP_k. All the results are presented in terms of applicability, i.e., the proportion of the instances to which AP is applicable, and are averaged over 20 runs.

	IP_1	IP_2	IP_3
AP_1	0.3800		
AP_2	0.8121	0.1456	
AP_3	0.7421	0.6590	0.1735
AP_4	0.6867	0.5667	0.5654

we tested $AP_{k+2}, ..., AP_4$ on IP_k to obtain their performances on IP_k, i.e., $P(AP_{k+2}, IP_k), ..., P(A_4, IP_k)$, and calculated the performance drops from $P(AP_{k+1}, IP_k)$ to these performances, i.e., $|P(AP_{k+1}, IP_k) - P(AP_{k+2}, IP_k)|$,..., $|P(AP_{k+1}, IP_k) - P(A_4, IP_k)|$, then these performance drops were compared to the performance improvement. The averaged performances (over 20 runs) of $AP_k, ..., AP_4$ on IP_k are presented in Table 5. The averaged performance improvement on IP_1 is $P(AP_2, IP_1) - P(AP_1, IP_1) = 0.4321$, and the two averaged performance drops on IP_1 are $P(AP_3, IP_1) - P(AP_2, IP_1) = 0.0700$ and $P(AP_4, IP_1) - P(AP_2, IP_1) = 0.1254$, so the ratios between the performance drops and the performance improvements on IP_1 are 16.20% and 29.02%. Calculated in the same way, the ratio on IP_2 is 17.98%. All the ratios between the performance drops and the corresponding performance improvements are below 30%.

In order to check whether the performances of AP on $IP_{training}$ are improved by LiangYi, the algorithm populations obtained from each iteration of LiangYi, i.e., AP_1, AP_2, AP_3, AP_4, were tested on $IP_{training}$. The testing results averaged over 20 runs are depicted in Figure 4. A constant improvement of the performances of AP_k on $IP_{training}$, according to the increase of k, is shown.

4.4.2. Testing Part

The algorithm population obtained from each iteration of LiangYi, i.e., $AP_1, ..., AP_4$ was tested on $IP_{testing}$. The testing results averaged over 20 runs are depicted in Figure 5. Once again, a constant improvement of the performances of AP_k on $IP_{testing}$ according to the increase of k is shown.

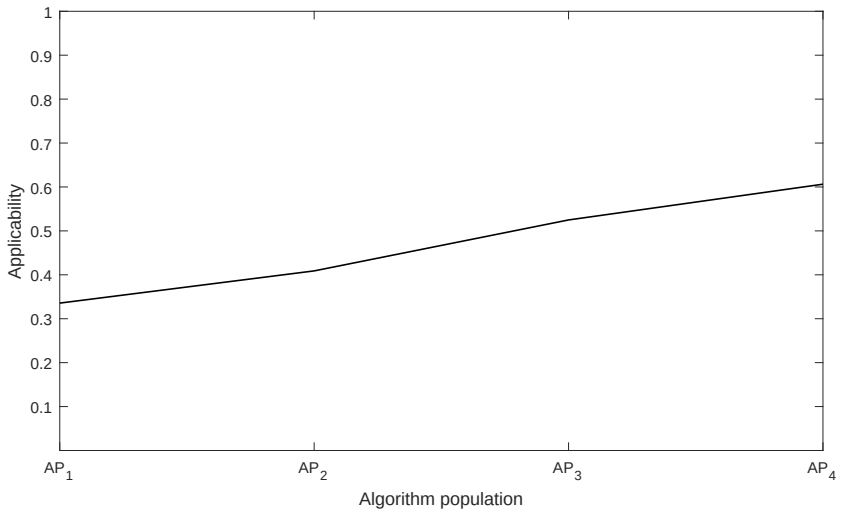

Figure 4: The performances of AP_1, AP_2, AP_3, AP_4 on $IP_{training}$. The performances are represented in terms of applicability, i.e., the proportion of the instances to which AP is applicable, and are averaged over 20 runs.

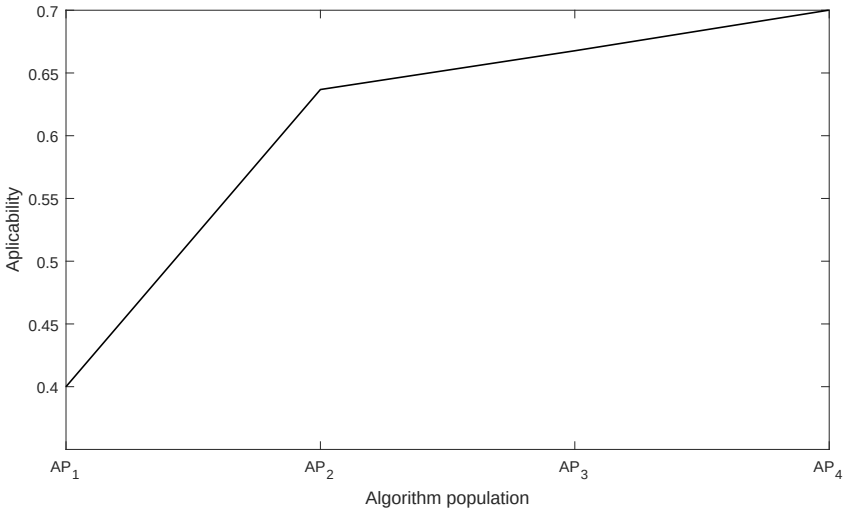

Figure 5: The performances of AP_1, AP_2, AP_3, AP_4 on $IP_{testing}$. The performances are represented in terms of applicability, i.e., the proportion of the instances to which AP is applicable, and are averaged over 20 runs.

5. Conclusion

This paper first put forward the concept of Experience-based optimizer (EBO) whose performance can improve as it solves more and more problem instances, and summarized several previous research in a unified context, i.e., offline training of EBO. A new training method, dubbed LiangYi, was proposed. The most novel feature of LiangYi is that, different from existing methods, it addresses selecting training instances and training solvers simultaneously. A specific instantiation of LiangYi on TSP was also proposed. Empirical results showed the advantages of LiangYi in comparison to ParHydra, the state-of-the-art APC method, on a huge testing set containing 10000 instances. Moreover, through empirically investigating behaviours of LiangYi, we confirmed that Liangyi is able to continuously improve the solver through training.

As discussed in the introduction, EBO is a far more broad direction than merely offline training of problem solvers. Further investigations may include:

(1) Further improvements to LiangYi. Diversity preservation scheme, such as speciation [**?**] or negatively correlated search [**?**] can be introduced into LiangYi to explicitly promote cooperation between different algorithms in AP. Another tack is to use machine learning techniques to accelerate LiangYi. Specifically, regression models and classification models can be used to predict the performances of algorithms in AP or instances in IP, without actually evaluating them, which is vary time-consuming.

(2) Online mode of EBO. Situations in which a solver faces a series of different problem instances coming sequentially pose new challenges. For example, the objective in online mode is to maximize the cumulative performances on all the instances. Thus methods designed for this scenario must consider making solvers perform well on current instances and improving solvers for future instances simultaneously. Besides, in a dynamic environment, the underlying properties of instances may change overtime; therefore the solvers may keep detecting the changes of environment and adapt to new instances.

(3) Deep understanding of the fundamental issues of EBO is also worthy of exploration. For example, LiangYi actually maintains two adversary sets competing with one another, which is a typical scenario where game theory can be applied. Besides, other more general issues in EBO include the similarity measure between instances, a unified approach to information extraction from solved instances, and theoretical proofs of the usefulness of transmitting information between similar instances.

6. Acknowledgements

References